UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF ILLINOIS
EASTERN DIVISION

CheckSum Ventures, LLC, Plaintiff, v. **Dell Inc.,** Defendant.	Case No. 1:18-cv-06321-RMD Patent Case Jury Trial Demanded

FIRST AMENDED COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff CheckSum Ventures LLC ("CheckSum"), through its attorney, complains of Dell, Inc. ("Dell"), and alleges the following:

PARTIES

1. Plaintiff CheckSum Ventures LLC is a corporation organized and existing under the laws of Texas that maintains its principal place of business at 3324 S Keaton Ave, Tyler, TX 75701.

2. Defendant Dell Inc. is a corporation organized and existing under the laws of Delaware that maintains its principal place of business at One Dell Way, Round Rock, TX 78662.

JURISDICTION

3. This is an action for patent infringement arising under the patent laws of the United States, Title 35 of the United States Code.

4. This Court has exclusive subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a).

5. This Court has personal jurisdiction over Dell because it has engaged in systematic and continuous business activities in this District. Specifically, Dell provides its full range of services to residents in this District. As described below, Dell has committed acts of patent infringement giving rise to this action within this District.

VENUE

6. Venue is proper in this District under 28 U.S.C. § 1400(b) because Dell has an established places of business in this district, including one located at 10 S Riverside Plaza, Chicago, IL 60606; and Dell has committed acts of patent infringement in this District. In addition, CheckSum has suffered harm in this district.

PATENT-IN-SUIT

7. CheckSum is the assignee of assignee of all right, title and interest in United States Patent No. 8,301,906 (the "'906 Patent" or the "Patent-in-Suit"), including all rights to enforce and prosecute actions for infringement and to collect damages for all relevant times against infringers of the Patent-in-Suit. Accordingly, CheckSum possesses the exclusive right and standing to prosecute the present action for infringement of the Patent-in-Suit by Dell.

The '906 Patent

8. On October 30, 2012, the United States Patent and Trademark Office issued the '906 Patent. The '906 Patent is titled "Apparatus for Writing Information on a Data Content on a Storage Medium." The application leading to the '906 Patent was filed on July 27, 2007 and is a national stage entry and continuation of the PCT application PCT/EP2007/003658 filed on April

25, 2007. A true and correct copy of the ’906 Patent is attached hereto as Exhibit A and incorporated herein by reference.

9. The ’906 Patent is valid and enforceable.

THE ’906 PATENT EMBODIES INVENTIVE CONCEPTS DESCRIBED IN THE SPECIFICATION AND CAPTURED IN THE CLAIMS

10. The claimed invention squarely addresses problems plaguing the prior art. Until the claimed invention entered the scene, data administration in the prior art lacked the technology to allow other users to securely verify the owner’s data:

The prior art did not allow other users to securely verify the owner’s data
“Conventional data administration concepts lack the possibility for users to allow other users to verify or integrity check data. Especially when using storage media that season and tend to become more and more erroneous with time it is a problem that at some point one can no longer be sure of the data validity or consistency, i.e. if the data can still be retrieved correctly from Such a medium.” (’906 patent, 1:27-34.)

11. Prior art storage systems also did not allow a user to securely *verify* and *prove* the origin and integrity of data:

The prior art did not allow users to securely verify and prove data origin
“Moreover conventional storage concepts and storage media do not allow [a user] to verify an origin of data. For example if data is transferred using portable storage media, e.g. by sending a CD (CD=Compact Disc) or a DVD (DVD-Digital Versatile Disk) by mail, the receiver cannot easily prove the origin of the data, i.e. verify the integrity of the data.” (*Id.*, 1:27-34.)

12. The specification also shows how to implement core facets of the claimed invention.

13. One inventive feature, the ability of another user to securely verify the integrity of an owner’s data content, is both claimed and extensively described in the specification. *See id.,* ’906 patent, 4:38-43 (“[A] user can use a private key to encrypt the checksums, another user can verify the checksums by decrypting them with a public key to obtain the decrypted checksum

information, which can be verified against checksum information obtained from the data content.”). This checksum information can be stored and tracked by referencing “the allocations through a pointer stored in a certain sector.” *Id.,* 2:61-63. And “assigning a particular checksum to its respective file can be done through a chunk table in an embodiment specifying a logical sector number of a first data block of a file and a checksum the file is associated with.” *Id.*, 2:64-67; *see also id.*, claim 6.

14. This claimed, particular use of checksum information, in conjunction with the claimed readers, is entirely integral to this claimed technique, and therefore not merely limiting a conventional activity to a special application of “providing checksum information” to different readers.

15. The claimed checksum information must be stored at a *particular* physical or logical location of the storage medium, this location must be tracked in the form of “control information,” and the enhanced reader (*not* the baseline reader) must be able to process and read the control information. *See, e.g.*, *id.* claim 1. The claims therefore provide a *particular inventive technique*—not for moving around data—but for allowing another user to securely verify the integrity of an owner’s data content.

16. The Patent-In-Suit also describes and claims specific technological improvements to address the limitations in early data security and content verification.

17. Specifically, and according to the specification, “[e]mbodiments of the present invention therefore provide the advantage that data can be verified, and a user can be prevented from working with broken data. Moreover, an effective mechanism is enabled to verify an origin of data stored on a storage medium. Some embodiments support public key signatures for optical storage media. Using this technology, the authenticity of a disc can be proven by verifying a

digital signature stored on the disc against a public verification key that needs to be provided once by an author of optical media. The digital signature refers to a checksum of the data on the storage medium. Some embodiments can use the private counterpart to the verification key to digitally design a hash value generated over the checksums." *Id.,* 2:39-51.

18. In addition to the data content and the checksum information, control information on the physical or logical location of the checksum information is written and used for verifying the checksum information. The enhanced reader can read and process the control information and the checksum information, while the prior art baseline readers could not do this and ignored, skipped or did not read the checksum information. Using the claimed invention allowed backwards compatibility with the prior art baseline readers, while still allowing enhanced readers to verify authenticity of encrypted checksums in accordance with the claimed invention. The prior art teaches away from verifying stored data, as message dependent checksum values calculated for verifying the accuracy of data in transit would only ensure the accuracy of transmitted information, and would not detect errors in the underlying stored data that may have been introduced in the storage medium itself.

19. Having such a system, combined with encrypted data verification techniques for the underlying stored data that is also backwards compatible with prior art baseline readers, is therefore both novel and inventive.

COUNT I: INFRINGEMENT OF THE '906 PATENT

20. CheckSum incorporates the above paragraphs herein by reference.

21. **Direct Infringement.** Dell has been and continues to directly infringe at least claim 1 of the '906 Patent in this District and elsewhere in the United States by providing products, for example, Dell's EMC VNX Storage Systems, which writes checksum information

(e.g., a SHA-1 hash value) for stored content on that storage system ("Dell's EMC VNX System"). The checksum is calculated and written for at least the purposes of deduplication.

22. Dell's EMC VNX System satisfies claim element 1(p) because it has: "[a]n apparatus for writing checksum information on a data content on a storage medium." For example, Dell's EMC VNX System writes checksum information (e.g., a SHA-1 hash value) for stored content on that storage system:

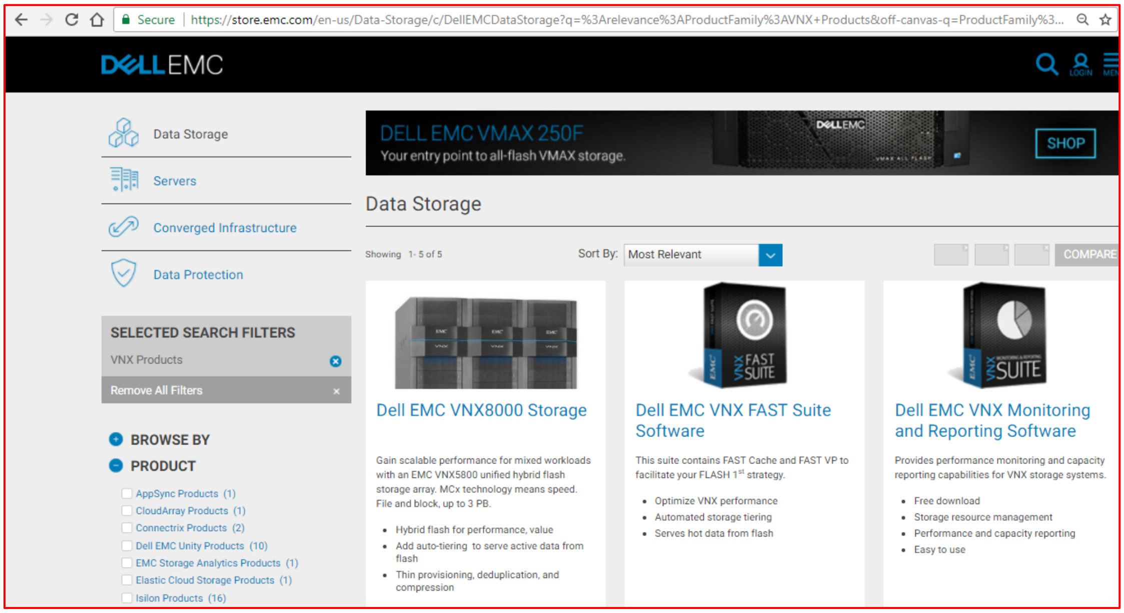

Secure | https://store.emc.com/en-us/Data-Storage/c/DellEMCDataStorage?q=%3Arelevance%3AProductFamily%3AVNX+Products&off-canvas-q=ProductFamily%3...

DELL EMC

Data Storage

Servers

Converged Infrastructure

Data Protection

DELL EMC VMAX 250F

Your entry point to all-flash VMAX storage.

Data Storage

Showing 1- 5 of 5

Sort By: Most Relevant

SELECTED SEARCH FILTERS

VNX Products

Remove All Filters ×

BROWSE BY

PRODUCT

- □ AppSync Products (1)
- □ CloudArray Products (1)
- □ Connectrix Products (2)
- □ Dell EMC Unity Products (10)
- □ EMC Storage Analytics Products (1)
- □ Elastic Cloud Storage Products (1)
- □ Isilon Products (16)

Dell EMC VNX8000 Storage

Gain scalable performance for mixed workloads with an EMC VNX5800 unified hybrid flash storage array. MCx technology means speed. File and block, up to 3 PB.

- Hybrid flash for performance, value
- Add auto-tiering to serve active data from flash
- Thin provisioning, deduplication, and compression

Dell EMC VNX FAST Suite Software

This suite contains FAST Cache and FAST VP to facilitate your FLASH 1st strategy.

- Optimize VNX performance
- Automated storage tiering
- Serves hot data from flash

Dell EMC VNX Monitoring and Reporting Software

Provides performance monitoring and capacity reporting capabilities for VNX storage systems.

- Free download
- Storage resource management
- Performance and capacity reporting
- Easy to use

Available at: https://store.emc.com/en-us/Data-Storage/c/DellEMCDataStorage?q=%3Arelevance%3AProductFamily%3AVNX+Products&off-canvas-q=ProductFamily%3AVNX+Products&facetselected=true; webpage attached hereto as Exhibit B.

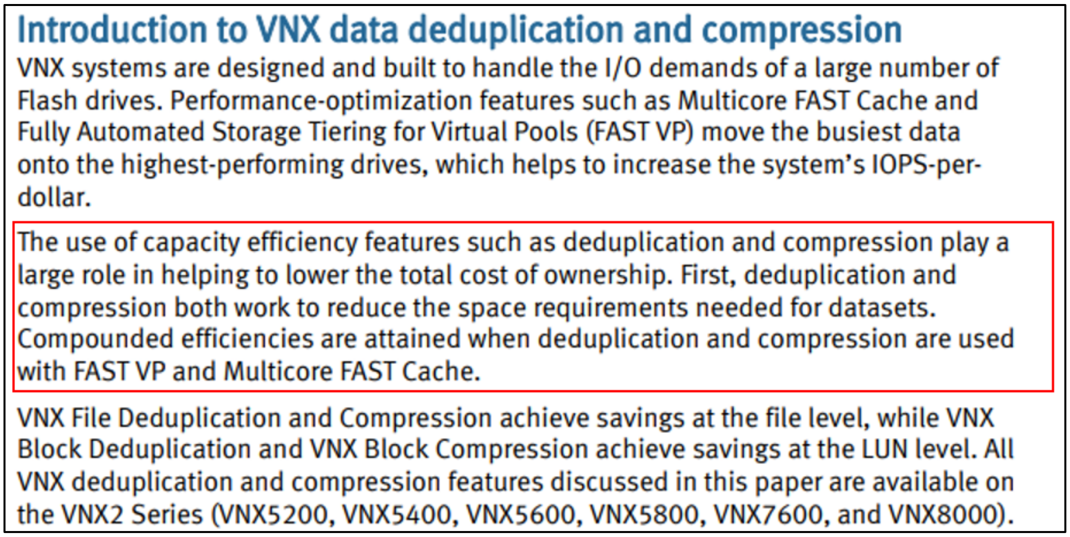

Introduction to VNX data deduplication and compression

VNX systems are designed and built to handle the I/O demands of a large number of Flash drives. Performance-optimization features such as Multicore FAST Cache and Fully Automated Storage Tiering for Virtual Pools (FAST VP) move the busiest data onto the highest-performing drives, which helps to increase the system's IOPS-per-dollar.

The use of capacity efficiency features such as deduplication and compression play a large role in helping to lower the total cost of ownership. First, deduplication and compression both work to reduce the space requirements needed for datasets. Compounded efficiencies are attained when deduplication and compression are used with FAST VP and Multicore FAST Cache.

VNX File Deduplication and Compression achieve savings at the file level, while VNX Block Deduplication and VNX Block Compression achieve savings at the LUN level. All VNX deduplication and compression features discussed in this paper are available on the VNX2 Series (VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, and VNX8000).

Available at: https://www.emc.com/collateral/white-papers/h12209-vnx-deduplication-compression-wp.pdf (p. 5); attached hereto as Exhibit C.

23. Dell's EMC VNX System satisfies claim element 1(a) because it has: "a provider for providing checksum information based on a data content." For example, Dell's EMC VNX System writes checksum information (e.g., a SHA-1 hash value) for stored content on that storage system. The checksum is calculated and written for at least the purposes of deduplication and is based on the data content:

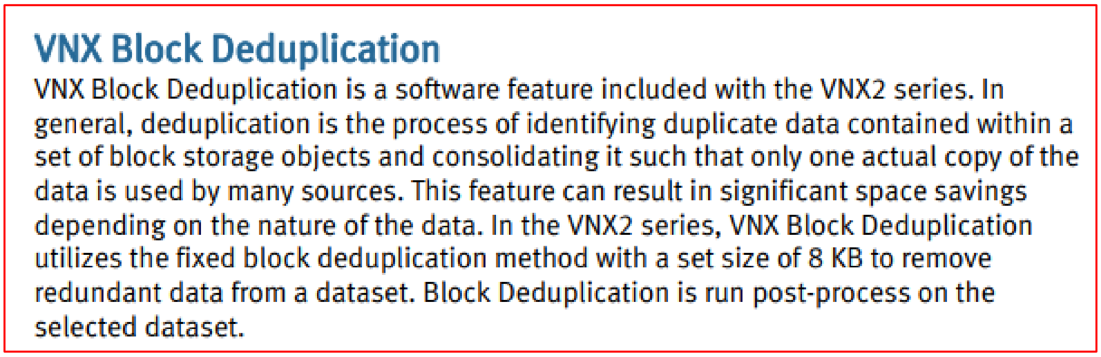

VNX Block Deduplication

VNX Block Deduplication is a software feature included with the VNX2 series. In general, deduplication is the process of identifying duplicate data contained within a set of block storage objects and consolidating it such that only one actual copy of the data is used by many sources. This feature can result in significant space savings depending on the nature of the data. In the VNX2 series, VNX Block Deduplication utilizes the fixed block deduplication method with a set size of 8 KB to remove redundant data from a dataset. Block Deduplication is run post-process on the selected dataset.

See Ex. C, p. 6.

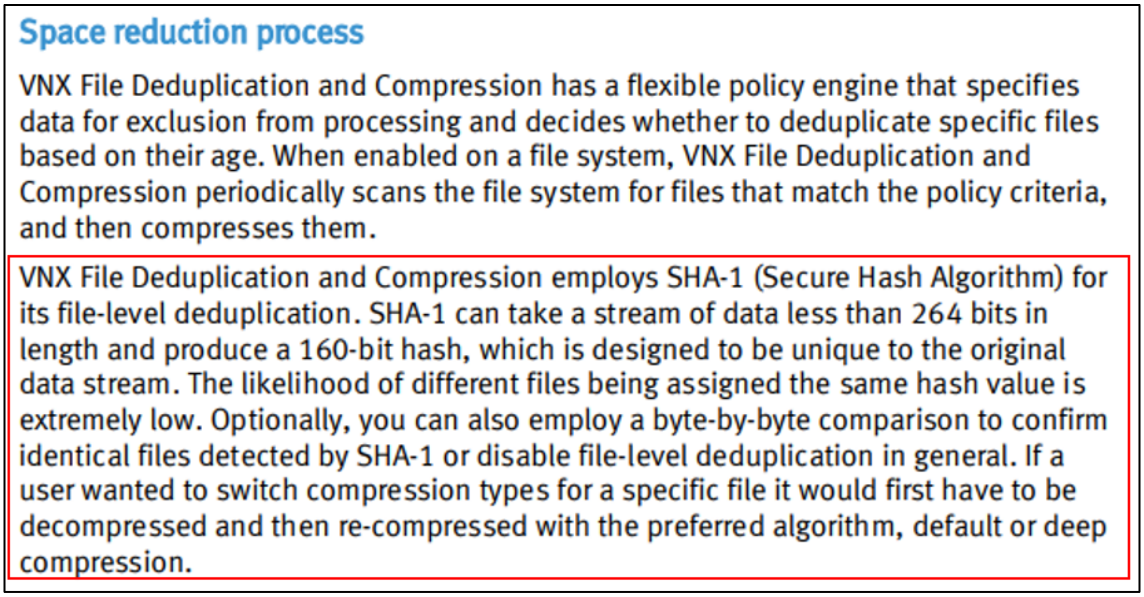

Space reduction process

VNX File Deduplication and Compression has a flexible policy engine that specifies data for exclusion from processing and decides whether to deduplicate specific files based on their age. When enabled on a file system, VNX File Deduplication and Compression periodically scans the file system for files that match the policy criteria, and then compresses them.

VNX File Deduplication and Compression employs SHA-1 (Secure Hash Algorithm) for its file-level deduplication. SHA-1 can take a stream of data less than 264 bits in length and produce a 160-bit hash, which is designed to be unique to the original data stream. The likelihood of different files being assigned the same hash value is extremely low. Optionally, you can also employ a byte-by-byte comparison to confirm identical files detected by SHA-1 or disable file-level deduplication in general. If a user wanted to switch compression types for a specific file it would first have to be decompressed and then re-compressed with the preferred algorithm, default or deep compression.

Available at: https://www.emc.com/collateral/hardware/white-papers/h8198-vnx-deduplication-compression-wp.pdf, p. 7; attached hereto as Exhibit D.

24. Dell's EMC VNX System satisfies claim element 1(b) because it has: "a writer for writing the data content, the checksum information and control information on a physical or logical location of the checksum information on the storage medium, such that a baseline reader and an enhanced reader can read the data content, the enhanced reader can read and process the control information and the checksum information and the baseline reader ignores, skips or does not read the checksum information." For example, Dell's EMC VNX System writes checksum information (e.g., a SHA-1 hash value) for stored content in a logical and/or physical location on that storage system.

25. Dell uses the stored checksum information at least in order to deduplicate the stored data. For example, when the same file is stored (or is being stored) in the storage system, Dell's EMC VNX System calculates one or more hash values for the file's contents. If the hash value is one that matches any of the hash values corresponding to files already stored on the

storage, EMC VNX deletes (or does not store) the second copy of the file – it simply stores an indication pointing to the logical and/or physical location of the original file on the storage:

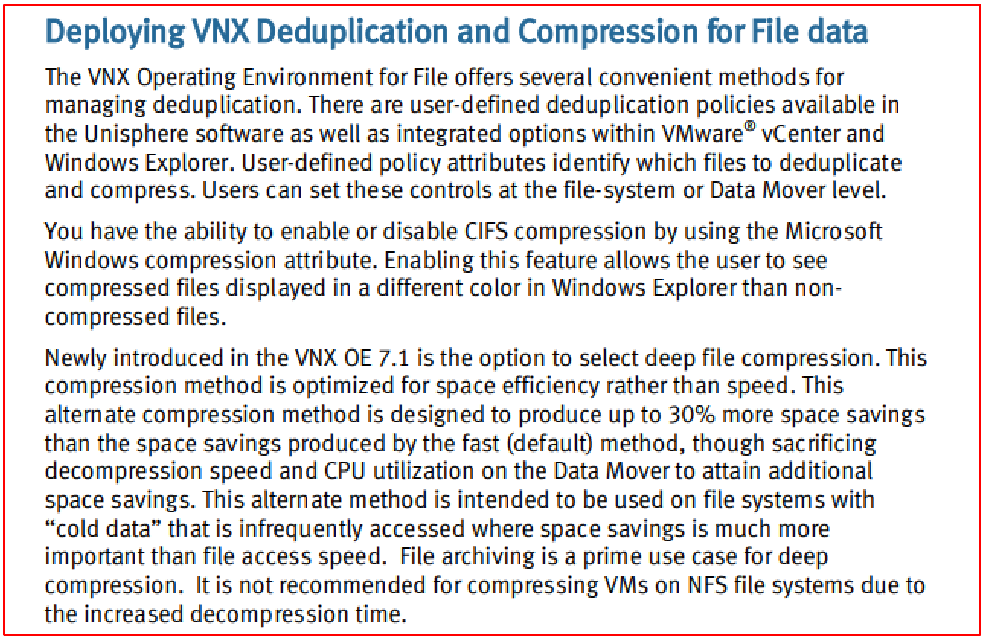

Deploying VNX Deduplication and Compression for File data

The VNX Operating Environment for File offers several convenient methods for managing deduplication. There are user-defined deduplication policies available in the Unisphere software as well as integrated options within VMware® vCenter and Windows Explorer. User-defined policy attributes identify which files to deduplicate and compress. Users can set these controls at the file-system or Data Mover level.

You have the ability to enable or disable CIFS compression by using the Microsoft Windows compression attribute. Enabling this feature allows the user to see compressed files displayed in a different color in Windows Explorer than non-compressed files.

Newly introduced in the VNX OE 7.1 is the option to select deep file compression. This compression method is optimized for space efficiency rather than speed. This alternate compression method is designed to produce up to 30% more space savings than the space savings produced by the fast (default) method, though sacrificing decompression speed and CPU utilization on the Data Mover to attain additional space savings. This alternate method is intended to be used on file systems with "cold data" that is infrequently accessed where space savings is much more important than file access speed. File archiving is a prime use case for deep compression. It is not recommended for compressing VMs on NFS file systems due to the increased decompression time.

See Ex. D, p. 11.

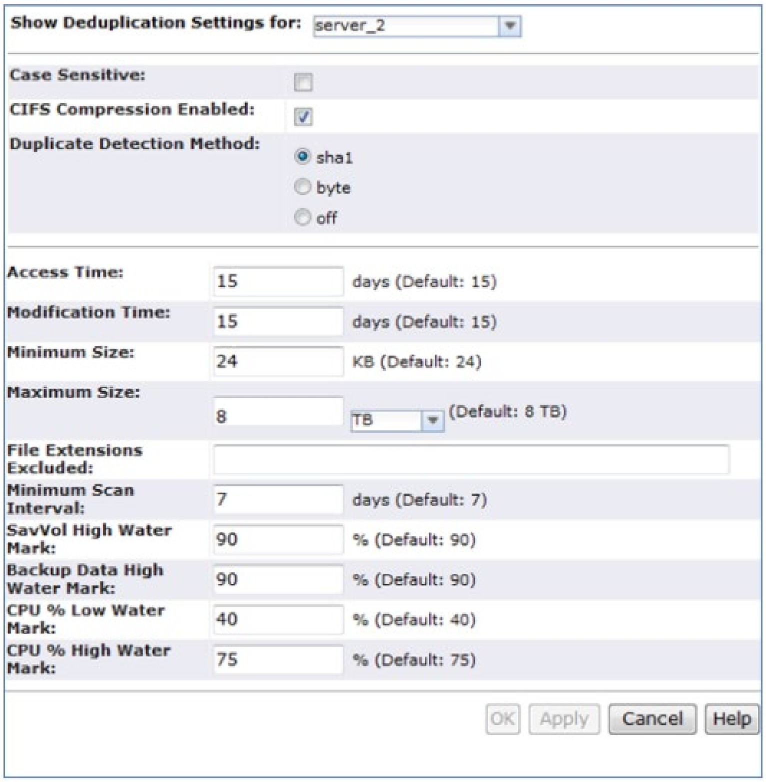

Show Deduplication Settings for: server_2

Case Sensitive:	□
CIFS Compression Enabled:	☑
Duplicate Detection Method:	● sha1 ○ byte ○ off

Access Time:	15	days (Default: 15)
Modification Time:	15	days (Default: 15)
Minimum Size:	24	KB (Default: 24)
Maximum Size:	8	TB (Default: 8 TB)
File Extensions Excluded:		
Minimum Scan Interval:	7	days (Default: 7)
SavVol High Water Mark:	90	% (Default: 90)
Backup Data High Water Mark:	90	% (Default: 90)
CPU % Low Water Mark:	40	% (Default: 40)
CPU % High Water Mark:	75	% (Default: 75)

Figure 3. Deduplication Settings at the Data Mover level

See Ex. D, p. 12.

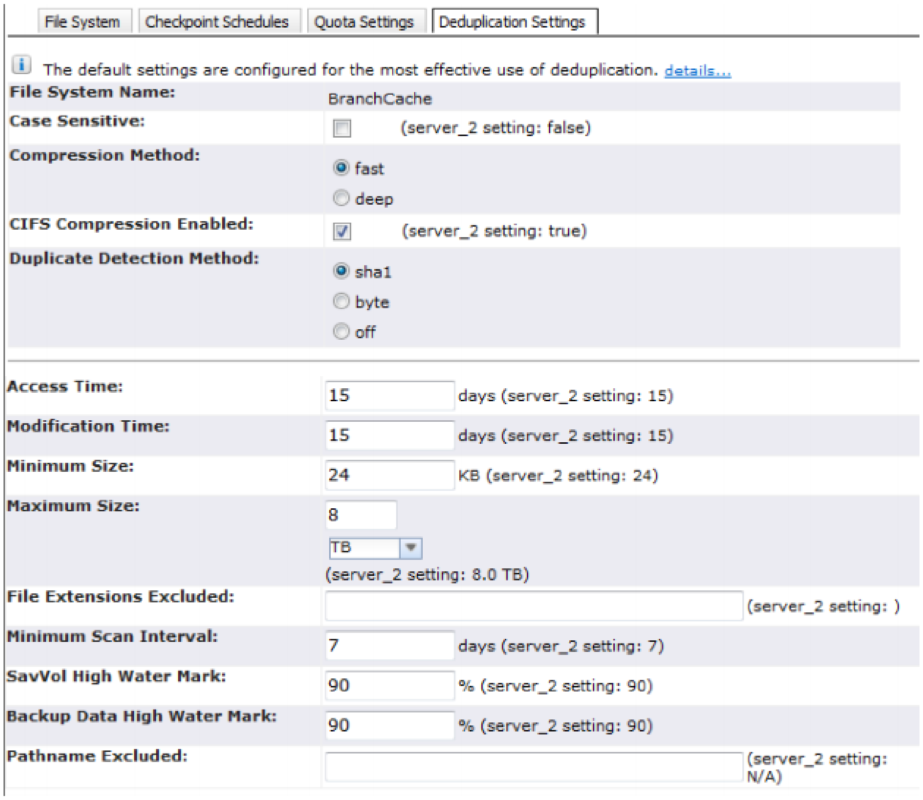

The default settings are configured for the most effective use of deduplication. details...

File System Name:	BranchCache
Case Sensitive:	□ (server_2 setting: false)
Compression Method:	● fast ○ deep
CIFS Compression Enabled:	☑ (server_2 setting: true)
Duplicate Detection Method:	● sha1 ○ byte ○ off
Access Time:	15 days (server_2 setting: 15)
Modification Time:	15 days (server_2 setting: 15)
Minimum Size:	24 KB (server_2 setting: 24)
Maximum Size:	8 TB (server_2 setting: 8.0 TB)
File Extensions Excluded:	(server_2 setting:)
Minimum Scan Interval:	7 days (server_2 setting: 7)
SavVol High Water Mark:	90 % (server_2 setting: 90)
Backup Data High Water Mark:	90 % (server_2 setting: 90)
Pathname Excluded:	(server_2 setting: N/A)

Figure 4. Deduplication Settings tab at the file system level

See Ex. D, p. 13.

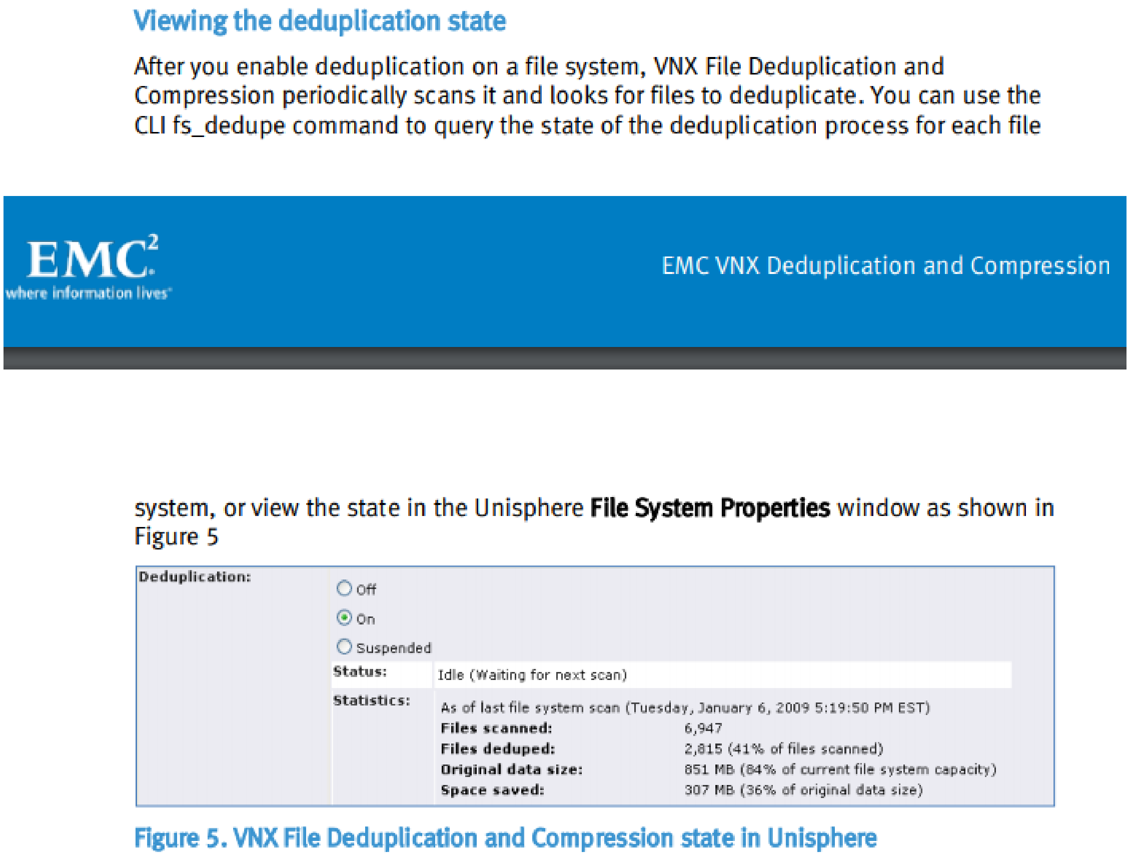

Viewing the deduplication state

After you enable deduplication on a file system, VNX File Deduplication and Compression periodically scans it and looks for files to deduplicate. You can use the CLI fs_dedupe command to query the state of the deduplication process for each file

EMC² where information lives®

EMC VNX Deduplication and Compression

system, or view the state in the Unisphere **File System Properties** window as shown in Figure 5

Deduplication:
○ Off
● On
○ Suspended

Status: Idle (Waiting for next scan)

Statistics: As of last file system scan (Tuesday, January 6, 2009 5:19:50 PM EST)
Files scanned: 6,947
Files deduped: 2,815 (41% of files scanned)
Original data size: 851 MB (84% of current file system capacity)
Space saved: 307 MB (36% of original data size)

Figure 5. VNX File Deduplication and Compression state in Unisphere

See Ex. D, pp. 13-14.

26. Dell's EMC VNX System also includes an enhanced reader that reads and processes the hash values such that, for example, the enhanced reader retrieves the original file even when the storage system queries the storage for the second file.

27. Dell's EMC VNX System also includes a Fast Cache and/or DRAM, into which the storage system stores data it recently retrieved. Whenever Dell's EMC VNX System performs a read operation, such as requesting a particular datum to be retrieved from the storage, the storage system first checks the Fast Cache and/or DRAM for the data. The storage system only checks for the datum if it is absent from the Fast Cache and/or DRAM. If it is absent from the Fast Cache and/or DRAM, it checks the main deduplicated storage for the requested data.

Thus, Dell's EMC VNX System also includes a baseline reader, including at least the reader operating on a FAST Cache and/or DRAM, which does not process the checksum information.

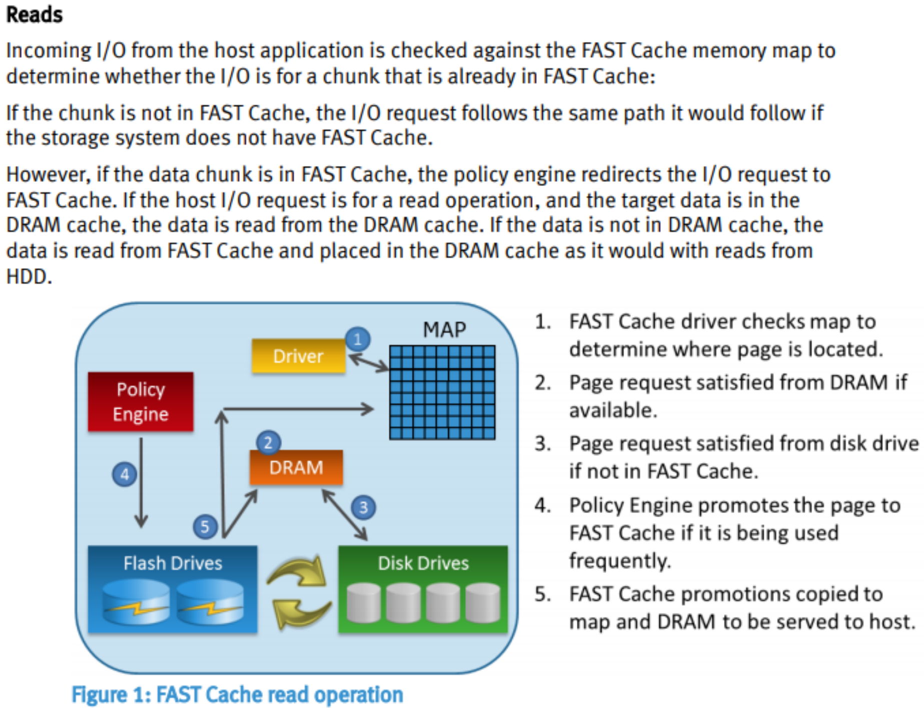

Reads

Incoming I/O from the host application is checked against the FAST Cache memory map to determine whether the I/O is for a chunk that is already in FAST Cache:

If the chunk is not in FAST Cache, the I/O request follows the same path it would follow if the storage system does not have FAST Cache.

However, if the data chunk is in FAST Cache, the policy engine redirects the I/O request to FAST Cache. If the host I/O request is for a read operation, and the target data is in the DRAM cache, the data is read from the DRAM cache. If the data is not in DRAM cache, the data is read from FAST Cache and placed in the DRAM cache as it would with reads from HDD.

1. FAST Cache driver checks map to determine where page is located.
2. Page request satisfied from DRAM if available.
3. Page request satisfied from disk drive if not in FAST Cache.
4. Policy Engine promotes the page to FAST Cache if it is being used frequently.
5. FAST Cache promotions copied to map and DRAM to be served to host.

Figure 1: FAST Cache read operation

Available at: https://www.emc.com/collateral/software/white-papers/h8046-clariion-celerra-unified-fast-cache-wp.pdf, (p. 9); attached hereto as Exhibit E.

28. CheckSum is entitled to recover damages adequate to compensate it for such infringement in an amount no less than a reasonable royalty under 35 U.S.C. § 284.

Jury Demand

29. Under Rule 38(b) of the Federal Rules of Civil Procedure, CheckSum respectfully requests a trial by jury on all issues so triable.

Prayer for Relief

WHEREFORE, CheckSum asks this Court to enter judgment against Dell, granting the following relief:

A. A declaration that Dell has infringed the Patent-in-Suit;

B. An award of damages to compensate CheckSum for Dell's direct infringement of the Patent-in-Suit;

C. An award of damages, including trebling of all damages, sufficient to remedy Dell's willful infringement of the Patent-in-Suit under 35 U.S.C. § 284;

D. A declaration that this case is exceptional, and an award to CheckSum of reasonable attorneys' fees, expenses and costs under 35 U.S.C. § 285;

E. An award of prejudgment and post-judgment interest; and

F. Such other relief as this Court or jury may deem proper and just.

Dated: October 28, 2019

Respectfully submitted,

/s/ Isaac Rabicoff
Isaac Rabicoff
Rabicoff Law LLC
73 W Monroe St
Chicago, IL 60603
(773) 669-4590
isaac@rabilaw.com

Counsel for Plaintiff

CERTIFICATE OF SERVICE

The undersigned certifies that a copy of the foregoing document was served on all parties who have appeared in this case on October 28, 2019, via the Court's CM/ECF system.

/s/ Isaac Rabicoff
Isaac Rabicoff